

A Refined Measurement of the Mean Transmitted Flux in the Ly α Forest over $2 < z < 5$ Using Composite Quasar Spectra

George D. Becker^{1*}, Paul C. Hewett¹, Gábor Worseck² and J. Xavier Prochaska²

¹*Kavli Institute for Cosmology and Institute of Astronomy, Madingley Rd, Cambridge, CB3 0HA, UK*

²*Department of Astronomy and Astrophysics, UCO/Lick Observatory, University of California, 1156 High Street, Santa Cruz, CA 95064, USA*

Draft version 14 August 2012

ABSTRACT

We present new measurements of the mean transmitted flux in the Ly α forest over $2 < z < 5$ made using 6065 quasar spectra from the Sloan Digital Sky Survey DR7. We exploit the lack of evolution of the mean quasar continuum to avoid the bias introduced by continuum fitting at high redshifts, which has been the primary systematic uncertainty in previous measurements of the mean Ly α transmission. The individual spectra are first combined into twenty-six composites with mean redshifts spanning $2.25 \leq z_{\text{comp}} \leq 5.08$. The flux ratios of separate composites at the same rest wavelength are then used, without continuum fitting, to infer the mean transmitted flux, $F(z)$, as a fraction of its value at $z \sim 2$. Absolute values for $F(z)$ are found by scaling our relative values to literature measurements made from high-resolution data at $z \leq 2.5$, where continuum uncertainties are minimal. We find that $F(z)$ evolves smoothly with redshift, with no evidence of a previously reported feature at $z \simeq 3.2$. This trend is consistent with a gradual evolution of the ionization and thermal state of the intergalactic medium over $2 < z < 5$. Our results generally agree with the most careful measurements to date made from high-resolution data, but offer much greater precision and extend to higher redshifts. This work also improves upon previous efforts using SDSS spectra, which have largely suffered from systematic errors. This is the most accurate and precise measurement of the mean transmitted flux at these redshifts to date, and should serve as a reference for future studies that use the Ly α forest.

Key words: intergalactic medium - quasars: absorption lines - cosmology: observations - cosmology: large-scale structure of the Universe

1 INTRODUCTION

The pattern of absorption lines imprinted on the spectra of distant objects by neutral hydrogen in the intergalactic medium (IGM), known as the “Ly α forest”, is one of the most fundamental probes of cosmic structure. The characteristics of this absorption reflect the density distribution, ionization state, and temperature of the intergalactic gas. The gas, in turn, closely traces the underlying distribution of dark matter, albeit with significant deviations arising from hydrodynamics coupled to the radiative and mechanical feedback from galaxies and AGN. Consequently, the Ly α forest allows us to probe the baryon physics connecting the

evolution of large-scale structure to the highly non-linear processes driving galaxy formation.

A key requirement for utilizing the Ly α forest to study either cosmology or baryon physics is to accurately establish the basic parameters describing the flux distribution. The most basic of these is the mean transmitted flux, $F = \langle e^{-\tau} \rangle$, and its evolution with redshift. The mean flux, which is often expressed in terms of an effective optical depth, $\tau_{\text{eff}}(z) = -\ln F(z)$, directly constrains, for example, the intensity of the metagalactic ionizing background (e.g., Rauch et al. 1997; McDonald & Miralda-Escudé 2001; Bolton et al. 2005). It also influences statistics such as the flux probability distribution function and power spectrum, among others, which are used to constrain quantities ranging from the temperature of the IGM (e.g., Schaye et al. 2000; Zaldarriaga et al. 2001; Lidz et al. 2010; Becker et al.

* gdb@ast.cam.ac.uk

2011) to the free-streaming length of dark matter particles (e.g., Viel et al. 2008). Ly α forest studies increasingly rely on comparing real data to artificial absorption spectra drawn from simulations, which must be calibrated to the correct mean flux. With the increasing availability of high-quality quasar and gamma ray burst spectra taken with large telescopes, and the large number of moderate-resolution quasar spectra gathered with surveys such as the Sloan Digital Sky Survey (SDSS; York et al. 2000) and the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2012), a precise measurement of $F(z)$ is therefore critical for accurately extracting the maximum amount of information from the forest.

The main challenge in measuring $F(z)$ is to overcome cosmic variance while accurately estimating the unabsorbed continuum of the background objects, which have so far typically been quasars. Studies using high-resolution (generally $R \sim 40,000$) spectra have the advantage that the peaks of the Ly α forest transmission should approach the continuum (e.g., Rauch et al. 1997; McDonald et al. 2000; Schaye et al. 2003; Songaila 2004; Kirkman et al. 2005; Kim et al. 2007; Becker et al. 2007; Faucher-Giguère et al. 2008; Dall’Aglio et al. 2008; Becker et al. 2011). Even at moderate redshifts ($z \sim 2 - 3$), however, voids in the IGM will have a non-negligible optical depth. The continuum will therefore lie somewhat above the tops of the transmission peaks, a bias that increases towards higher redshifts. Alternatively, studies using large sets of moderate-resolution spectra have either attempted to simultaneously solve for $F(z)$ and the quasar continua in a strictly statistical sense (Bernardi et al. 2003), or have fit continua on an individual basis using principle component analysis (McDonald et al. 2005; Pâris et al. 2011) or with bias corrections estimated from model spectra (Dall’Aglio et al. 2009). At the highest redshifts ($z > 5$), $F(z)$ measurements have generally relied on extrapolating a power-law continuum from redward of the Ly α emission line (Songaila 2004; Fan et al. 2006; Becker et al. 2007).

The most accurate $F(z)$ measurement to date has arguably been made by Faucher-Giguère et al. (2008). These authors used a set of 86 high-resolution quasar spectra and made statistical corrections to the continua as a function of redshift based on artificial spectra drawn from hydrodynamic simulations. Notably, they identified a possible sharp ($\Delta z \lesssim 0.5$) dip in the evolution of the mean opacity near $z \simeq 3.2$. A similar feature was found by Bernardi et al. (2003) using a very different data set and measurement technique. The optical depth of the optically thin IGM to Ly α will scale with the H I neutral fraction, which in turn depends on the gas temperature, T , and the H I ionization rate, Γ , as $f_{\text{H I}} \propto T^{-0.7} \Gamma^{-1}$, where the temperature dependence is for case A recombination. Bernardi et al. noted that a decrease in $\tau_{\text{eff}}(z)$ at $z \simeq 3.2$ could indicate a temporary increase in the IGM gas temperature accompanying the end of helium reionization (see also Theuns et al. 2002; Faucher-Giguère et al. 2008). However, the reality of this feature remains unclear. Although a similar dip was found at modest significance by Dall’Aglio et al. (2008) in a sample of 40 high-resolution spectra, neither McDonald et al. (2005) nor Dall’Aglio et al. (2009) detected such a feature using samples of moderate-resolution SDSS spectra larger than that of Bernardi et al. (2003). Indeed,

Faucher-Giguère et al. (2008) emphasized that the significance of their feature may have been exaggerated by the assumption of an underlying power law evolution in τ_{eff} . The association of such a sharp feature with helium reionization has also been challenged on theoretical grounds. Bolton et al. (2009) argued that even if the IGM is photoheated during the course of a very rapid ($\Delta z \simeq 0.2$) helium reionization, the expected signature would be a sudden decrease in τ_{eff} followed by a gradual recovery driven by adiabatic cooling in the voids. Moreover, recent temperature measurements indicate that heating due to helium reionization was an extended process with $\Delta z \gtrsim 1$ (Becker et al. 2011), and so it is unclear why such a sharp feature in the opacity evolution should occur.

In this paper we present a new and highly precise measurement of the mean transmitted Ly α flux over $2 < z < 5$ using a sample of 6065 moderate-resolution quasar spectra drawn from the seventh data release of the Sloan Digital Sky Survey (SDSS DR7). The main innovation of this work is the introduction of a new method that uses composite spectra to measure the overall shape the mean flux evolution to high accuracy without fitting continua. These results are then normalized to measurements made from high-resolution data at $z \sim 2$, where continuum errors are minimal. Our approach builds on the successful use of composite quasar spectra to statistically constrain other properties of the IGM, for example the mean free path of ionizing photons (Prochaska et al. 2009). We also take advantage of the improved statistical accuracy afforded by DR7 to increase substantially the number of quasars analyzed compared to previous works (Bernardi et al. 2003; McDonald et al. 2005; Dall’Aglio et al. 2009; Pâris et al. 2011, although Pâris et al. also used DR7 data). This paper should therefore serve as a reference for future studies that utilize the mean Ly α opacity. In Section 2 we describe our technique, which uses composite spectra to compute $F(z)$ relative to that at $z \sim 2$. Our results are presented in Section 3. In Section 4, we compare our measurements to results from the literature. Finally, we summarize in Section 5.

2 DATA & METHOD

The normalized transmitted flux in a single spectrum is given by the observed flux divided by the unabsorbed continuum, $F^i = f_{\text{obs}}^i / f_{\text{cont}}^i$. The goal of this paper is to measure the mean normalized flux as a function of observed wavelength (i.e., redshift) across our entire sample of objects, $F(z) = \langle F^i(z) \rangle$, where $1 + z = \lambda_{\text{obs}} / \lambda_{\text{Ly}\alpha}$, and $\lambda_{\text{Ly}\alpha} = 1215.67 \text{ \AA}$. Rather than fit continua to individual spectra, however, we use the fact that the mean continuum for a large group of quasars will be highly similar regardless of redshift. The difference between composites at different redshifts should therefore primarily reflect the evolution in the mean transmitted flux in the Ly α forest. In this section we describe how our composites are constructed, and how these are used to calculate $F(z)$.

2.1 Constructing composite spectra

We constructed 26 composite quasar spectra using a total of 6065 individual spectra from the SDSS DR7 quasar cata-

Table 1. Composite spectra used in this work. Columns give the minimum and maximum quasar redshifts, the mean composite redshift, and the number of quasars included in each composite.

z_{\min}	z_{\max}	z_{comp}	\mathcal{N}
2.2	2.3	2.250	530
2.3	2.4	2.346	442
2.4	2.5	2.450	439
2.5	2.6	2.545	379
2.6	2.7	2.648	277
2.7	2.8	2.750	191
2.8	2.9	2.852	233
2.9	3.0	2.951	334
3.0	3.1	3.049	346
3.1	3.2	3.151	374
3.2	3.3	3.249	330
3.3	3.4	3.347	280
3.4	3.5	3.449	195
3.5	3.6	3.554	185
3.6	3.7	3.650	245
3.7	3.8	3.749	238
3.8	3.9	3.847	217
3.9	4.0	3.955	144
4.0	4.1	4.047	150
4.1	4.2	4.145	115
4.2	4.3	4.248	97
4.3	4.4	4.346	70
4.4	4.6	4.495	101
4.6	4.8	4.694	63
4.8	5.0	4.896	54
5.0	5.2	5.083	36

logue (Schneider et al. 2010). The redshift range, mean redshift, and the number of quasars contributing to each composite are listed in Table 1. Redshift bins of $\Delta z = 0.1$ were used between $z = 2.2$ and 4.4 . Above $z = 4.4$ we used a bin size of $\Delta z = 0.2$ in order to increase the number of objects in each composite. The quasars were selected to have an absolute magnitude $M_i > -29.0$, while lacking broad absorption line (BAL) features. For quasars present in the SDSS DR6 release, BAL quasars were defined using the classifications of Allen et al. (2011), while the small fraction of quasars present only in DR7 were classified via visual inspection. None of the results presented here, however, are sensitive to the exact definition of the BAL quasar sample. We also excluded objects where the peak of the Ly α emission line was greater than ten times the continuum level near a rest wavelength of 1280 Å. This was done primarily to avoid very high-redshift objects with extremely low signal-to-noise ratios in the continuum ($S/N \lesssim 2$). Improved quasar redshifts¹ were determined using the Hewett & Wild (2010) scheme for generating self-consistent redshifts for quasars with SDSS spectra. We used spectra that had been corrected for telluric emission line residuals in the red (Wild & Hewett 2005, 2010). In addition, for roughly 6 per cent of our objects, spectra from multiple epochs were combined to increase the signal-to-noise ratio.

The SDSS spectra are supplied on a logarithmic wavelength scale, with pixels of 69 km s^{-1} width. When generating composite spectra in the quasar rest-frame, nearest-pixel values were employed without interpolation or rebinning.

¹ http://www.sdss.org/dr7/products/value_added/#quasars




Figure 1. Continuum power law slopes for the quasars used to construct our composites, where $f_\lambda \propto \lambda^\alpha$, and α is measured between 1270 Å and 1510 Å (rest wavelengths). Black points give the values for individual objects. Red points with error bars show the mean value and standard deviation in each redshift bin. The horizontal line at $\alpha = -1.56$ shows the composite value from Vanden Berk et al. (2001).

The arithmetic mean over all contributing objects was then computed at each pixel. All spectra were weighted equally in order to minimize the effects of cosmic variance on the flux in the Ly α forest. Error arrays were also computed by directly computing the intra-sample standard deviation. For the initial composites (described below), pixels falling within 16 Å intervals centered on the wavelengths of strong [O I] sky emission lines at 5578.5 and 6301.7 Å were masked out. For the final composites, we masked all regions contaminated by strong telluric lines that are flagged in the bitmask accompanying each SDSS spectrum. We note that the flagged regions typically do not include all of the pixels in a spectrum affected by skylines, and often cover only the core region contaminated by a skyline and not the line wings, which can still be significantly affected by residuals. Using spectra that have been corrected for skyline residuals therefore still proved beneficial.

Initial composites were generated using the original quasar spectra without modification, except to divide each object by its mean flux near 1285 Å, which is the region of the continuum closest to the Ly α forest that is relatively free of emission lines. These composites exhibited an encouraging similarity in both continuum and emission-line properties longward of 1250 Å in the quasar rest-frame. A weak but systematic trend, however, appeared in the overall spectral shape as a function of redshift. To correct for this, we re-calculated the composites after adjusting the continuum slope of each individual object using the following procedure. The high signal-to-noise ratio initial composite at $z \simeq 2.65$ was chosen as a reference. For each individual quasar, a power law was then fit to the ratio of the flux in its spectrum to that in the reference composites, and the spectrum was divided by this fit at all wavelengths. The fit was made between rest-frame wavelengths 1270 Å and 1510 Å, which is the reddest rest-frame wavelength region covered in the SDSS spectra for objects at $z \sim 5$. Note that we avoid the high-ionization lines Ly α λ 1216, N V λ 1239,1243, and C IV λ 1548,1551, which tend to vary strongly between




Figure 2. The twenty-six composite quasar spectra used in this work, plotted versus rest-frame wavelengths. The spectra are binned using 345 km s^{-1} pixels, as used in the analysis, and separated into multiple panels for clarity. The mean quasar redshift is also displayed for each composite. The vertical dotted lines indicate the region of the Ly α forest used to measure $F(z)$. Note that the prominence of the emission lines in the forest naturally decreases with increasing composite redshift as the overall transmitted flux decreases.

objects. We also exclude the region between 1380 \AA and 1410 \AA , which contains the Si IV $\lambda 1394, 1403$ emission line doublet.

The best-fitting spectral indices, calculated by taking the slope of the initial $z \simeq 2.65$ composite as a reference point, are plotted in Figure 1. We note that higher-redshift objects tend to be systematically redder, albeit slightly. This trend may reflect a combination of effects related both to the selection of the quasars in the SDSS and intrinsic changes in the quasar spectra, or even aspects of the data reduction. It is clearly essential, however, to remove any systematic trend that could impact the determination of the measurement of the Ly α -forest depression with redshift. The only assumption underlying the nature of the correction procedure adopted is that the essentially featureless changes to the continuum slopes, which were measured over large wavelength intervals redward of the Ly α emission, extend into the spectral region used to study the Ly α forest. We test this assumption in Section 3.

The final composites generated from the continuum-corrected SDSS spectra are shown in Figure 2. The composite spectra are strikingly similar redward of the Ly α +N V emission lines. The overall amplitudes and slopes are expected to be the same based on the continuum correction procedure described above. We note, however, that the amplitudes of the emission lines redward of Ly α +N V are also very similar. This is particularly true of the low ionization

lines, O I $\lambda 1302$ + Si II $\lambda 1304$ and C II $\lambda 1334$, which is encouraging for our purposes since the lines in the forest are also from low ionization species (e.g. Vanden Berk et al. 2001).

Significant variations do appear in the strength of the Ly α emission line, which, taken at face value, might be of concern. The effect, however, is due to errors in the redshifts of the individual objects, which, because of Ly α absorption on the blue side of the Ly α emission line, tend to decrease the height of the emission in the composite spectrum. The effect increases towards higher redshift as the absorption in the Ly α forest increases. The redshift uncertainties for the SDSS quasar spectra also increase with redshift (Table 3 of Hewett & Wild 2010) as the C IV emission line has to be used in the redshift determination. At the very highest redshift, $z_{\text{comp}} \simeq 5.08$, the effect apparently diminishes (see Figure 2) because the Ly α emission line itself dominates in the redshift determination, causing the peak of the composite Ly α line to sharpen by selection. By comparison, the flux over the region spanning $1250 \text{ \AA} < \lambda < 1500 \text{ \AA}$, which contains relatively weak emission lines, shows very little variation with redshift. Since the emission lines in the forest are also weak, redshift errors should have a minimal effect there, as well.

We directly verified that redshift uncertainties affect the strength of Ly α by comparing composite spectra of 56 quasars at $4.4 < z < 5.2$ with improved redshifts from




Figure 3. A demonstration of the effects of redshift errors on composite spectra. The same set of 56 $4.4 < z < 5.2$ quasar spectra were used to create both versions of the composite plotted here. The red line shows the composite as generated using individual redshifts measured from the SDSS spectra themselves, while for the black line the redshifts were measured from higher-quality GMOS spectra. The improved redshifts from the GMOS data produce a more sharply peaked Ly α emission line but have little effect on the rest of the spectrum. Note that, since we selected only objects where the GMOS redshifts could be measured from O I $\lambda 1302$ + Si II $\lambda 1304$ and/or C II $\lambda 1334$, the emission lines are somewhat stronger than the composites shown in Figure 2.

recently-obtained Gemini/GMOS spectra (Worseck et al, in prep). Two composites were generated using the same set of SDSS spectra. In one case we use the redshifts measured directly from the SDSS spectra, while in the other case we use redshifts from O I $\lambda 1302$ + Si II $\lambda 1304$ and/or C II $\lambda 1334$ measured from the GMOS data. The results are shown in Figure 3. The improved GMOS redshifts produce a stronger and more sharply peaked Ly α line, as expected, while the remainder of the spectrum is nearly unaffected. This supports our assumption that, for all of our composites, the differences in the fluxes blueward of the Ly α emission lines are due solely to the evolution of the mean opacity of the Ly α forest. It is these differences that we will use to constrain $F(z)$.

2.2 Calculating the mean transmitted flux

At a given rest wavelength λ_0 , the flux of a quasar composite with mean emission redshift z_{comp} will be attenuated by a factor

$$F(z) = \frac{f_{\text{obs}}(\lambda_0, z_{\text{comp}})}{f_{\text{cont}}(\lambda_0)}, \quad (1)$$

where $1 + z = \lambda_0(1 + z_{\text{comp}})/\lambda_{\text{Ly}\alpha}$. We unfortunately do not have a completely unabsorbed composite at $z_{\text{comp}} = 0$ that would allow us to measure $f_{\text{cont}}(\lambda_0)$ absolutely, and hence calculate $F(z)$ directly. However, we can measure the ratio of $F(z)$ at two redshifts by evaluating the ratio of fluxes between two composites at the same rest-frame wavelengths,

$$\frac{F(z_1)}{F(z_2)} = \frac{f_{\text{obs}}(\lambda_0, z_{\text{comp},1})}{f_{\text{obs}}(\lambda_0, z_{\text{comp},2})}. \quad (2)$$

Here, the dependence on the continuum has been eliminated because we are using composites rather than individual spectra. This relation will hold to the degree that the mean continuum in the composites either does not evolve or has been adequately corrected with redshift, an assumption we test in Section 3. Nominally, however, this dataset allows us to compute $F(z)$ up to a normalization factor by finding a function that satisfies the set of observed flux ratios.

We parametrized $F(z)$ as a stepwise function in bins of $\Delta z = 0.1$, with the first bin centered on $z = 2.15$. We then computed $F(z)$ as a fraction of the transmitted flux in that bin, i.e., $F(z)/F(z = 2.15)$. When analyzing the data, we bin the composite fluxes into 345 km s^{-1} pixels. This is to avoid large non-symmetric errors in $F(z_1)/F(z_2)$ when $F(z_2)$ becomes small, although nearly identical results were obtained without binning. We also made secondary adjustments to the composite continua by dividing each composite by a power-law fit over $1270 \text{ \AA} < \lambda < 1510 \text{ \AA}$ to the ratio of that composite to the average of all composites with $2.25 < z_{\text{comp}} < 4.25$. This was a minor correction to help remove any residual differences in the composite shapes, and changed the final $F(z)$ values by less than the 1σ errors.

The primary goal of this paper is to measure the mean opacity of intergalactic hydrogen to Ly α . Here we loosely define “intergalactic” to include all diffuse, highly-ionized absorbers with H I column densities $N_{\text{HI}} < 10^{19} \text{ cm}^{-2}$. As defined in equation 1, however, $F(z)$ in the Ly α forest will also include absorption from metal lines, as well as from higher-column density Ly α absorbers. We now describe how we account for these additional contributions to arrive at the desired $F(z)$ values.

2.2.1 Metal lines

Absorption in the Ly α forest due to intervening metals, while small compared to the Ly α absorption, must be accounted for in order to measure $F(z)$ precisely. Conventional studies using continuum-normalized spectra have relied on masking (e.g., Schaye et al. 2003) or fitting out (Kim et al. 2007) metal lines individually, or else applying global statistical corrections (e.g., Kirkman et al. 2005). Measuring flux ratios in composite spectra, however, offers an alternative approach to mitigating the impact of metals, as discussed below.

The mean level of metal-line absorption may vary, in principle, with both quasar redshift and rest-frame wavelength. A large fraction of absorption in the forest will come from the C IV $\lambda 1548, 1551$, Si IV $\lambda 1394, 1403$, and Mg II $\lambda 2796, 2804$ doublets, which, because they have rest wavelengths longward of Ly α , will produce, on average, a continuum of absorption extending over the region redward of the Ly α emission line we use to scale our spectra and into the forest. An additional contribution will come from the various metal lines associated with high H I column density systems such as DLAs. These lines have a variety of rest wavelengths, including some that fall between Ly β and Ly α .

In practice, the flux decrement from metals is relatively minor, and appears to depend only mildly on either quasar redshift or rest-frame wavelength. Becker et al. (2011) found a nearly constant decrement of ~ 3 per cent in the region between the N V $\lambda 1239, 1243$ and Si IV $\lambda 1394, 1403$ emission lines in the high-resolution spectra of a large sample

of quasars with redshifts $z \sim 2 - 4$. Kirkman et al. (2005) also found that the absorption due to metals redward of Ly α evolves slowly with quasar redshift. In the Ly α forest, Schaye et al. (2003) found a similar (~ 3 percent) level of absorption from metals over Ly α redshifts $z \sim 2 - 4$, with little evidence for strong evolution. Kim et al. (2007), in possibly the most careful such work to date, also found a mean decrement of ~ 3 per cent due to all metals in the forest over $z \sim 2 - 3$, albeit with significant scatter between sight lines, and no evidence for evolution.

In a composite spectrum, the absorption due to metals will essentially scale the mean flux by a multiplicative factor, and may also introduce a change in slope with rest-frame wavelength. Given the apparently mild dependence of this absorption on either rest-frame wavelength or quasar redshift, however, we expect these effects to be minor (on the order of a few percent). More importantly, the process of scaling the composites to have the same amplitude and slope between the Ly α and C IV emission lines will tend to eliminate the relative differences in the metal line absorption as a function of composite redshift. This is another advantage of our differential approach to measuring $F(z)$. Systematic errors with redshift may remain if the metal-line absorption (or more precisely, the relative difference in metal-line absorption between composites at different redshifts) is not well represented by a power law in wavelength (see Section 2.1). Given the metal-line measurements discussed above, however, we expect such residual systematics to be minor, even compared to relatively small statistical errors on $F(z)$ from such a large data set.

We note that establishing an absolute scaling for $F(z)$ still requires a correction for metal lines, at least over a narrow range in redshift. As described below, we use absolute values for $F(z \sim 2)$ that have been corrected following Schaye et al. (2003). This correction for metal lines is roughly consistent with other works, although it is still a potential source of systematic error in our final $F(z)$ values, albeit at the $\lesssim 1$ per cent level.

2.2.2 Correction for optically thick absorbers

The total Ly α opacity will include a contribution from absorbers that are optically thick to ionizing photons (i.e., with $N_{\text{HI}} \geq 10^{17.2} \text{ cm}^{-2}$). As discussed below, most of this absorption will come from super Lyman limit systems (SLLSs; $10^{19.0} \text{ cm}^{-2} \leq N_{\text{HI}} < 10^{20.3} \text{ cm}^{-2}$) and damped Ly α absorbers (DLAs; $N_{\text{HI}} \geq 10^{20.3} \text{ cm}^{-2}$). These are believed to be associated with intervening galaxies, and are difficult to model in numerical simulations. In order to facilitate a direct comparison with simulations, therefore, we correct for these absorbers in order to arrive at our final $F(z)$ values. Following Faucher-Giguère et al. (2008), we use a model to calculate the integrated absorption due to optically thick systems. For a given frequency distribution in column density, N_{HI} , Doppler parameter, b , and redshift, the flux decrement due to these absorbers will be

$$1 - F_{\text{T}}(z) = \frac{1+z}{\lambda_{\text{Ly}\alpha}} \int_{N_{\text{min}}}^{N_{\text{max}}} dN_{\text{HI}} \int db f(N_{\text{HI}}, b, z) W_0(N_{\text{HI}}, b), \quad (3)$$

where W_0 is the rest-frame equivalent width. At high column densities, W_0 will depend mainly on the column density, and so we set $b = 20 \text{ km s}^{-1}$, as in Faucher-Giguère et al.

(2008). We assume that the shape of the column density distribution does not evolve with redshift, such that $f(N_{\text{HI}}, z) = f(N_{\text{HI}}) dn/dz$. We use the broken power-law column density distribution at $z \sim 3.7$ compiled by Prochaska et al. (2010), and adopt 10^{19} cm^{-2} and 10^{22} cm^{-2} as our limits of integration. Lyman limit systems with $10^{17.2} \text{ cm}^{-2} \leq N_{\text{HI}} < 10^{19.0} \text{ cm}^{-2}$ were found to contribute a negligible amount of Ly α absorption at our level of precision, while DLAs with $N_{\text{HI}} > 10^{22} \text{ cm}^{-2}$ are vanishingly rare. The redshift evolution of the number density of SLLSs and DLAs is not known precisely over the entire redshift range; however, $dn/dz \propto (1+z)^2$ gives a reasonable fit to the evolution of DLAs over $2.3 < z < 4.4$ measured by Prochaska & Wolfe (2009). Adopting this evolution gives an integrated flux decrement due to SLLSs and DLAs of $1 - F_{\text{T}}(z) = 0.0045 [(1+z)/3]^{3.0}$. We note that this is somewhat more absorption than estimated by Faucher-Giguère et al., who adopted a model with $f(N_{\text{HI}}) \propto N_{\text{HI}}^{-1.5}$ over all column densities, and $dn/dz \propto (1+z)^{1.5}$.

When calculating $F(z)$, we first fit raw values relative to $z \sim 2$ that include absorption from both the optically thin Ly α forest and from optically thick absorbers. We then corrected these relative values using the above scaling for F_{T} to give the relative mean flux in the forest only. These corrected values of $F(z)/F(z = 2.15)$ were then scaled to measurements at $z \leq 2.5$ from the literature to derive absolute values for $F(z)$, as described below.

2.2.3 Absolute scaling

In order to compute absolute values for $F(z)$, we scaled our results to $F(z)$ measurements made from high-resolution spectra at redshifts where continuum uncertainties are minimal. We used the results of Faucher-Giguère et al. (2008), who used artificial spectra drawn from cosmological simulations to correct for continuum-fitting biases. These authors found the corrections to be small ($\lesssim 2$ per cent) at $z < 2.5$. Their results were also confirmed by Becker et al. (2011), who measured $F(z)$ from high-resolution spectra using a different approach to correct for continuum errors. Becker et al. re-normalized both the real data and the simulated spectra by deliberately placing the continuum along the tops of the highest transmission peaks. The true $F(z)$ was then measured from the simulations once the simulated optical depths had been adjusted such that the re-normalized data and simulations had the same mean flux. The resulting $F(z)$ agreed with the Faucher-Giguère et al. measurements to within one per cent over $2 < z < 2.5$. We note that Faucher-Giguère et al. and Becker et al. also used different methods to correct for metal-line contamination; the Faucher-Giguère et al. values adopted here include a correction based on metal-line measurements in the forest from Schaye et al. (2003), whereas Becker et al. applied a similar correction based on absorption redward of Ly α .

To scale our measurements, we interpolated our $F(z)/F(z = 2.15)$ results onto the centers of the four $\Delta z = 0.1$ redshift bins between $z = 2.2$ and 2.5 used by Faucher-Giguère et al. (2008). The mean flux ratio between our raw results and the Faucher-Giguère et al. values, weighted by the errors, was then used to compute our final $F(z)$. We note that when computing the scaling we first adjusted the Faucher-Giguère et al. data to account for the dif-

ference in the amount of absorption we assume for optically thick absorbers. This is a minor (~ 0.1 per cent) correction at $z \leq 2.5$, however.

2.2.4 Error estimates

Uncertainties in $F(z)$ were estimated using a bootstrap re-sampling approach. In each iteration, a new composite was generated for each redshift bin from a set of \mathcal{N} quasars drawn randomly, with replacement, from the set of \mathcal{N} objects in that bin. $F(z)$ was then re-computed from the new composites following the procedure described in Section 2.1. This process was repeated 1000 times, and the full covariance arrays for both the unscaled $F(z)/F(z = 2.15)$ and scaled $F(z)$ were computed from the results. We stress that, because we are fundamentally fitting ratios of fluxes at difference redshifts, the errors in $F(z)$ will be correlated. Any analysis that uses the $F(z)$ results should therefore consider the entire covariance array.

2.2.5 Impact of sky-subtraction errors

Measuring $F(z)$ at $z < 2.3$ requires analyzing fluxes at very blue wavelengths ($\lambda_{\text{obs}} < 4000$ Å). Recent papers using SDSS data (Prochaska et al. 2010; Pâris et al. 2011), however, have noted potential problems with the sky-subtraction in SDSS spectra at these wavelengths. Pâris et al., for example, found a net positive residual below 4000 Å in a sample of quasar spectra that should show no flux below 4280 Å due to the presence of a DLA at $z > 3.7$. In principle, a positive residual would cause us to overestimate $F(z)$ at $z < 2.3$, and would also impact our higher-redshift measurements, since the errors in our $F(z)$ points will be correlated. We note that, for our measurements, the impact of sky-subtraction residuals in the far blue is likely to be small; we only use data at $\lambda_{\text{obs}} < 4000$ Å in the lowest-redshift composites ($z_{\text{comp}} \leq 2.75$), and these were constructed using relatively bright quasars. Nevertheless, we performed multiple tests to determine whether sky-subtraction errors could significantly affect our $F(z)$ results.

We first directly estimated the mean sky residual (and any contamination from intervening objects within the SDSS fibers) by measuring the flux over $\lambda_{\text{obs}} = 3900 - 4000$ Å in the spectra of each of our quasars at $z > 4.5$. These objects should generally not have any flux below 4000 Å due to the integrated Lyman-series opacity of the IGM and any optically thick absorbers. A few exceptionally clear lines of sight may show some transmission, in which case the flux will be an upper limit on the sky-subtraction residuals. We measured a mean flux of $f_{4000} \simeq (8 \pm 1) \times 10^{-19} \text{ ergs}^{-1} \text{ cm}^{-2} \text{ Å}^{-1}$. For our quasars at $z < 2.8$, this is, on average, roughly 0.5 per cent of the observed flux over the same wavelength range. This is comparable to (albeit slightly larger than) than the 1σ errors on $F(z)/F(z = 2.15)$ at $z < 2.3$, as described below, and substantially less than the errors in the final scaled $F(z)$ values. This supports the conclusion that, since the quasars we use to measure the flux below 4000 Å are bright, sky-subtraction errors at these wavelengths will have a relatively minor impact.

Second, we tried adjusting the redshift interval over which we scaled our absolute $F(z)$ values. If sky-subtraction

errors are artificially increasing our $F(z)/F(z = 2.15)$ values at $z < 2.3$, the matching our results to literature values at these redshifts should yield a normalization factor that is too low. To test this, we tried scaling our results to the Faucher-Giguère et al. (2008) measurements over $2.6 \leq z \leq 3.0$, rather than over $2.2 \leq z \leq 2.5$. The resulting scale factor was higher by only 0.5 per cent, which is well within the 1σ uncertainty of ~ 1 per cent (see measurements below).

Finally, we tried fitting $F(z)$ using only data at absorption redshifts $z > 2.5$, or $\lambda_{\text{obs}} > 4250$ Å. This had a negligible impact on the shape of $F(z)$ at $z > 2.5$. In summary, therefore, sky-subtraction residuals in the far blue do not appear to have a significant impact on our results. We emphasize, however, that systematics such as these may need to be more carefully addressed in order to achieve substantially higher-precision mean flux measurements than the ones presented here.

3 RESULTS

Our results for $F(z)/F(z = 2.15)$, along with our $F(z)$ values scaled to the Faucher-Giguère et al. (2008) measurements at $z \leq 2.5$, are shown in Figure 4 and listed in Table 2. The errors are 1σ and reflect only the diagonal elements of the covariance matrix. For $F(z)$ the errors include the uncertainty in the overall absolute scaling, which has a fractional uncertainty of ~ 1 per cent (for the adopted set of corrections for continuum bias, metal-line contamination, and other systematics made by Faucher-Giguère et al.) and dominates the error budget at $z \leq 3.75$.

The mean flux exhibits a smooth decline with redshift over $2 < z < 5$. Notably, we see no indication of the feature at $z \sim 3.2$ reported by Bernardi et al. (2003) and Faucher-Giguère et al. (2008). Comparisons with these and other literature values are discussed in Section 4.

Our primary systematic uncertainty is whether the shape and amplitude of the mean unabsorbed continuum in the Ly α forest is the same for all composites, regardless of redshift. To check this, we first reconstructed the continuum in the forest for each composite, scaled by $F(z = 2.15)$, by dividing the observed fluxes by our raw (uncorrected for absorption due to optically thick systems) relative transmitted flux values,

$$f_{\text{cont}}(\lambda_0, z_{\text{comp}})F(z = 2.15) = \frac{f_{\text{obs}}(\lambda_0, z_{\text{comp}})}{F(z)/F(z = 2.15)}. \quad (4)$$

Here, z is calculated from the observed wavelength of each pixel, and the transmitted flux is interpolated between the values plotted in Figure 4. Observed wavelengths for a composite at redshift z_{comp} are related to the rest-frame wavelengths plotted in Figure 2 simply by $\lambda_{\text{obs}} = \lambda_0(1 + z_{\text{comp}})$. We then computed the mean scaled continuum across all composites with $z_{\text{comp}} < 4.05$, where the redshift limit was chosen to avoid noisier composites at higher redshifts. The result is shown in Figure 5. Finally, a model for the observed flux for each composite was created by multiplying the mean scaled continuum, $f_{\text{cont}}(\lambda_0)F(z = 2.15)$, by the relative mean transmitted flux, $F(z)/F(z = 2.15)$, evaluated separately for each pixel.

The fractional residuals from these models, $(f_{\text{obs}} - f_{\text{model}})/f_{\text{model}}$, are shown in Figure 6. Error bars were com-




Figure 4. The mean transmitted flux in the Ly α forest. Left panel: $F(z)$ as a fraction of the mean transmitted flux at $z = 2.15$. These values are computed directly from our composite spectra, and include a correction for the absorption due to optically thick systems. Right panel: $F(z)$ values after scaling our results to the measurements of Faucher-Giguère et al. (2008) over $2.2 \leq z \leq 2.5$. Error bars are 1σ , and include uncertainties in the scaling factor for $F(z)$ points in the right-hand panel. Tabulated values are given in Table 2. The full covariance matrices can be found in the accompanying supplemental material.




Figure 5. The mean quasar continuum in the Ly α forest. The continuum was calculated by averaging the composites with $z_{\text{comp}} < 4.05$ after correcting each one for the measured relative mean transmitted flux, $F(z)/F(z = 2.15)$, calculated separately at each observed wavelength. The overall amplitude of the continuum has therefore implicitly been multiplied by $F(z = 2.15)$. The solid line gives the results from the full sample. The shaded region shows the 1σ uncertainties calculated from bootstrap trials.

puted by repeating the above procedure for each bootstrap trial. The absence of significant features in the residuals indicates that the unabsorbed continuum is highly similar between composites. This includes the shape and amplitude of the emission lines as well as the overall slope of the continuum, which would introduce a tilt in the residuals if it varied

significantly between composites. The self-similarity of the continuum supports our assumption that the ratios of the observed fluxes between composites can be modeled purely by the evolution in $F(z)$ (plus a minor contribution from optically thick absorbers), at least up to a single normalization factor.

We emphasize that this test addresses the self-similarity of the continuum only over the Ly α forest. Variations with redshift in the observed mean quasar spectral energy distribution (SED) at other wavelengths are known to exist. For example, Worseck & Prochaska (2011) found that SDSS, in order to avoid contamination from the stellar locus, preferentially selects quasars with red SEDs in the u and g bands over $2.7 \lesssim z \lesssim 3.5$. They found that this primarily effects the observed mean flux for higher-order Lyman-series transitions, and has relatively little impact on the Ly α forest (see their Figure 17). Nevertheless, careful modeling of such selection effects may be required in order to achieve significantly more precise mean flux measurements than the ones presented here.

3.1 Analytic function

In addition to fitting a discrete $F(z)$, we fit an analytic function to the effective optical depth as a function of redshift, $\tau_{\text{eff}}(z) = -\ln F(z)$. It has been conventional to parametrize τ_{eff} as a simple power law of the form $\tau_{\text{eff}}(z) = \tau_0 [(1+z)/(1+z_0)]^\beta$. If one assumes this form for $\tau_{\text{eff}}(z)$, then both the normalization and the slope will be directly constrained from flux ratios, since neither τ_0 nor β cancels when computing $F(z_1)/F(z_2)$. With the precision and long redshift baseline provided by our dataset, however, we found




Figure 6. Fractional flux residuals in the Ly α forest for each of the twenty-six composites, plotted with respect to the corresponding Ly α absorption redshifts. The residuals are calculated as $(f_{\text{obs}} - f_{\text{model}})/f_{\text{model}}$, where f_{model} is the product of the mean continuum shown in Figure 5 and the relative mean transmitted flux, $F(z)/F(z = 2.15)$, evaluated at each observed wavelength (i.e., redshift). 1σ error bars were computed from bootstrap trials. A unity offset is added between composites for clarity. The lack of significant residuals demonstrates that the continuum over the Ly α forest (i.e., the slope of the underlying power-law continuum and the relative shape and amplitude of the emission lines) changes little with quasar redshift, and that the ratios of observed fluxes between composites can be well modeled by the evolution in the mean Ly α forest opacity.




Figure 7. Mean transmitted flux in the Ly α forest expressed as an effective optical depth, $\tau_{\text{eff}}(z) = -\ln F(z)$. The solid line is an analytic fit to the full data set (not the binned points displayed), using a modified power law of the form $\tau_{\text{eff}}(z) = \tau_0 [(1+z)/(1+z_0)]^\beta + C$. The shaded region gives the 1σ uncertainty in the model τ_{eff} values, taking into account the covariance between the model parameters. See Section 3.1 for details. The axes are scaled such that a pure power law with $C = 0$ would appear as a straight line.

that, although a pure power law could provide a reasonable fit to the observed flux ratios (i.e., produce an $F(z)$ with the correct shape), the overall normalization of $F(z)$ was substantially too low when compared to literature values. We found, however, that a reasonable fit could be achieved by allowing for a constant offset to $\tau_{\text{eff}}(z)$, such that

$$\tau_{\text{eff}}(z) = \tau_0 \left(\frac{1+z}{1+z_0} \right)^\beta + C. \quad (5)$$

Following our approach for the discrete $F(z)$, we fit τ_0 and β directly using the flux ratios, then solved for C by scaling our results to those of Faucher-Giguère et al. (2008) over $2.2 \leq z \leq 2.5$. Bootstrapping was again used to compute error estimates, including the covariance between fit parameters.

The results for $\tau_{\text{eff}}(z)$ are plotted in Figure 7. The best-fitting parameters are $[\tau_0, \beta, C] = [0.751, 2.90, -0.132]$, for $z_0 = 3.5$, and the full covariance matrix is given in Table 3. We emphasize that τ_0 and β were constrained directly from the composite flux ratios (corrected for absorption due to optically thick systems) rather than the binned $F(z)$ values. Nevertheless, the analytic fit closely matches the discrete values at all redshifts once both are scaled to the literature values at $z \sim 2$. In $F(z)$ space, the deviations of the discrete $F(z)$ from the nominal analytic fit have a formal χ^2 of 64, calculated using the full covariance matrix, for 28 redshift bins. This suggests that some of the residual structure in the discrete fit may be real, that a modified power law is not a sufficiently accurate model for $\tau_{\text{eff}}(z)$, or that we have somewhat underestimated the errors in the discrete case. Even so, the departures of the discrete points from the analytic fit are small, with a mean fractional deviation in the flux of two per cent. A modified power law for $\tau_{\text{eff}}(z)$ thus appears to be a

reasonable approximation to the true evolution of the Ly α opacity over this redshift range, although we caution that the fit should not necessarily be extrapolated to higher or lower redshifts.

4 COMPARISON TO PREVIOUS RESULTS

We now provide a brief comparison to some previous results from the literature. The most careful measurements of the mean Ly α flux using high-resolution spectra has arguably been made by Faucher-Giguère et al. (2008), whose values at $z \leq 2.5$ we have used to normalize our own results. Their values are plotted along with our results in Figure 8. Although we normalized our values to theirs only over $2.2 \leq z \leq 2.5$, the overall agreement out to $z = 4.2$ is very good. We note that at $z \geq 3.5$ the Faucher-Giguère et al. values for $\tau_{\text{eff}}(z)$ are systematically somewhat higher ($\Delta\tau_{\text{eff}} \sim 0.05$) than ours. This may indicate that the continuum corrections made by those authors were somewhat too large at the high-redshift end. More significantly, the feature in $\tau_{\text{eff}}(z)$ at $z \simeq 3.2$ tentatively identified by Faucher-Giguère et al. is absent from our results. In light of the substantially higher precision offered by the present analysis, the feature appears likely to have been spurious. It may either reflect random fluctuations in the Faucher-Giguère et al. data, or may be the result of a subtle systematic effect.

A potential concern is whether an analysis using composite spectra, which average the flux from objects over a range of redshifts, may smooth out discrete features in the mean flux evolution. We tested this by generating artificial composites with the same redshift range and binning as used

Table 2. Tabulated values for the mean transmitted flux in the Ly α forest. Columns give the redshift of each bin, the mean transmitted flux as a fraction of its value at $z = 2.15$ before and after correcting for absorption due to optically-thick absorbers, and the corrected absolute mean transmitted flux after scaling to the measurements of Faucher-Giguère et al. (2008) over $2.2 \leq z \leq 2.5$. The full covariance matrices can be found in the accompanying supplemental material.

z	$F(z)/F(z = 2.15)$		$F(z)$
	Raw	Corrected	
2.15	1.0000	1.0000	0.8806 ± 0.0103
2.25	0.9750 ± 0.0035	0.9755 ± 0.0035	0.8590 ± 0.0098
2.35	0.9421 ± 0.0044	0.9431 ± 0.0044	0.8304 ± 0.0093
2.45	0.9034 ± 0.0046	0.9049 ± 0.0046	0.7968 ± 0.0089
2.55	0.8849 ± 0.0049	0.8869 ± 0.0049	0.7810 ± 0.0090
2.65	0.8543 ± 0.0045	0.8568 ± 0.0045	0.7545 ± 0.0088
2.75	0.8341 ± 0.0044	0.8371 ± 0.0044	0.7371 ± 0.0088
2.85	0.8104 ± 0.0044	0.8139 ± 0.0044	0.7167 ± 0.0086
2.95	0.7871 ± 0.0045	0.7911 ± 0.0045	0.6966 ± 0.0084
3.05	0.7530 ± 0.0045	0.7574 ± 0.0045	0.6670 ± 0.0082
3.15	0.7203 ± 0.0044	0.7252 ± 0.0044	0.6385 ± 0.0080
3.25	0.6797 ± 0.0047	0.6849 ± 0.0047	0.6031 ± 0.0079
3.35	0.6487 ± 0.0047	0.6543 ± 0.0047	0.5762 ± 0.0074
3.45	0.6241 ± 0.0045	0.6301 ± 0.0045	0.5548 ± 0.0071
3.55	0.5983 ± 0.0046	0.6047 ± 0.0047	0.5325 ± 0.0071
3.65	0.5604 ± 0.0050	0.5669 ± 0.0050	0.4992 ± 0.0069
3.75	0.5295 ± 0.0053	0.5363 ± 0.0053	0.4723 ± 0.0068
3.85	0.5006 ± 0.0062	0.5076 ± 0.0062	0.4470 ± 0.0072
3.95	0.4759 ± 0.0061	0.4832 ± 0.0062	0.4255 ± 0.0071
4.05	0.4502 ± 0.0064	0.4576 ± 0.0065	0.4030 ± 0.0071
4.15	0.4177 ± 0.0069	0.4252 ± 0.0070	0.3744 ± 0.0074
4.25	0.4003 ± 0.0071	0.4081 ± 0.0073	0.3593 ± 0.0075
4.35	0.3828 ± 0.0105	0.3908 ± 0.0107	0.3441 ± 0.0102
4.45	0.3572 ± 0.0097	0.3652 ± 0.0099	0.3216 ± 0.0094
4.55	0.3337 ± 0.0110	0.3417 ± 0.0112	0.3009 ± 0.0104
4.65	0.3190 ± 0.0125	0.3272 ± 0.0129	0.2881 ± 0.0117
4.75	0.2674 ± 0.0219	0.2747 ± 0.0225	0.2419 ± 0.0201
4.85	0.2456 ± 0.0164	0.2527 ± 0.0168	0.2225 ± 0.0151

Table 3. Covariance matrix for $\tau_{\text{eff}}(z)$ parameters (equation 5). The best-fitting values are $[\tau_0, \beta, C] = [0.751, 2.90, -0.132]$, for $z_0 = 3.5$.

	τ_0	β	C
τ_0	0.00049	-0.00241	-0.00043
β	-0.00241	0.01336	0.00224
C	-0.00043	0.00224	0.00049

for the real data, i.e., bins of $\Delta z = 0.1$ from $z = 2.2$ to 4.4, and $\Delta z = 0.4$ from $z = 4.4$ to 5.2. Each composite was generated from 100 individual artificial spectra with redshifts distributed uniformly across the bin. For simplicity, a uniform continuum was assumed for each spectrum. The input Ly α forest absorption for each individual spectrum was calculated from the Faucher-Giguère et al. (2008) power law+Gaussian fit to $\tau_{\text{eff}}(z)$, which includes a Gaussian dip of width $\sigma_z \simeq 0.1$ centered at $z \simeq 3.2^2$. This is comparable to the feature reported by Bernardi et al. (2003). After adding noise to the artificial composites at a level comparable to that in the real data, we re-measured the mean flux evolu-

² See values using the Schaye et al. (2003) metal correction in Table 5 of Faucher-Giguère et al. (2008).




Figure 8. Our results as shown in Figure 7 plotted along with those of Faucher-Giguère et al. (2008), which are the most careful measurements to date made with high-resolution spectra. Although our values were scaled to the Faucher-Giguère et al. measurements only at $z \leq 2.5$, the overall agreement out to $z = 4.2$ is very good. In light of the precision offered by the new measurements presented here, the feature at $z \sim 3.2$ tentatively identified by Faucher-Giguère et al. appears to be consistent with noise. We note that the Faucher-Giguère et al. points over $2.2 \leq z \leq 2.5$, to which we normalized our results, appear noisy because they are presented on a log plot; their errors on the mean flux over this range are, in fact, relatively small ($\sigma_F \lesssim 0.02$).




Figure 9. A demonstration of the ability of our approach to recover a feature in the $\tau_{\text{eff}}(z)$ evolution similar to the one reported by Faucher-Giguère et al. (2008) and Bernardi et al. (2003). The solid line represents the input $\tau_{\text{eff}}(z)$, while the dotted line shows a pure power law for comparison. Points with error bars are the $\tau_{\text{eff}}(z)$ we recover by applying our analysis technique to artificial composite quasar spectra generated with the input $\tau_{\text{eff}}(z)$ but with properties otherwise similar to those in our sample. See Section 4 for details.

tion using the approach outlined in Section 2.2. The results, plotted in Figure 9, show that the input flux evolution was recovered with high accuracy. Although this was a simplified test, it demonstrates the key point that measurements of the mean flux evolution using composite spectra are entirely sensitive to discrete features similar to the ones that have been reported in the literature.

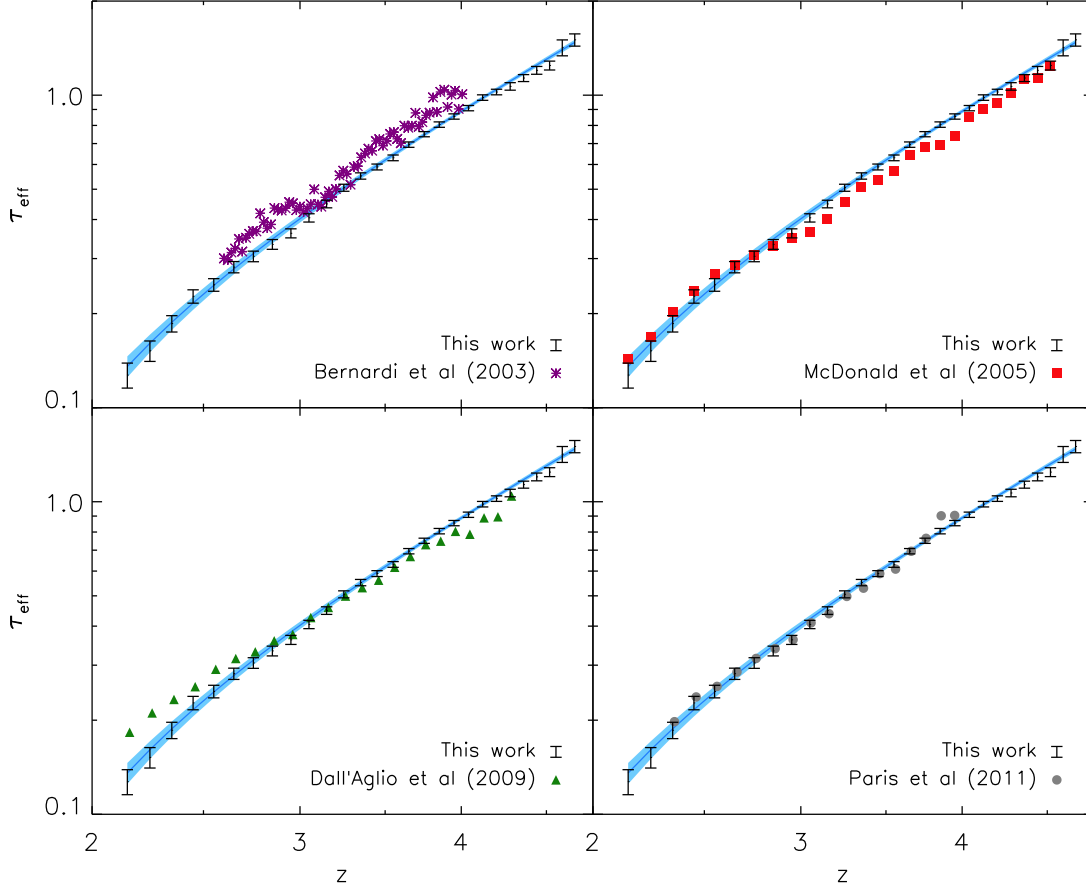


Figure 10. Our results as shown in Figure 7 plotted along with previous measurements made from SDSS spectra. Error bars have been omitted from the literature values for clarity and to highlight systematic trends. Deviations in the first three sets of measurements (Bernardi et al. 2003; McDonald et al. 2005; Dall’Aglio et al. 2009) from the results presented here are likely to reflect systematic errors in those works. The overall agreement between the Pâris et al. (2011) results and ours is generally good, although the Pâris et al. measurements have not been corrected for metal-line contamination. See Section 4 for discussion.

We next turn to previous measurements made using SDSS spectra. Results from Bernardi et al. (2003), McDonald et al. (2005), Dall’Aglio et al. (2009), and Pâris et al. (2011) are plotted along with our measurements in Figure 10. In the three earliest works there are systematic differences from our results. We propose that these differences are most likely to be explained by systematic errors in the previous analyses, as discussed below. In contrast, the agreement between the Pâris et al. (2011) measurements and our results is generally very good, which suggests that the measurements from SDSS spectra have now largely converged.

Bernardi et al. (2003) used composites generated from 1061 individual quasar spectra to simultaneously constrain $\tau_{\text{eff}}(z)$ and the mean unabsorbed continuum over the Ly α forest using two approaches. In the first approach, they conducted a joint fit to analytic functions for both $\tau_{\text{eff}}(z)$ and the mean continuum. Second, they used an iterative approach to fit a discrete $\tau_{\text{eff}}(z)$ along with an arbitrarily shaped continuum for each composite. Both approaches, however, face significant limitations. The composite spectra used by Bernardi et al., as well as the ones used here, fundamentally do not contain enough information to constrain

both the amplitude and shape of $F(z)$ (or $\tau_{\text{eff}}(z)$) unless additional assumptions are made. In their analytic approach, Bernardi et al. assumed a power law + Gaussian form for $\tau_{\text{eff}}(z)$ similar to the one used by Faucher-Giguère et al. (2008). As discussed in Section 3.1, the ratios of fluxes at different redshifts directly provided by composite spectra will constrain both the normalization and slope of a pure power law. In this case, therefore, $\tau_{\text{eff}}(z)$ will be implicitly normalized purely by the choice of fitting function (an additional, narrow Gaussian component will not change this since the contribution to $\tau_{\text{eff}}(z)$ will go to zero far from the Gaussian centroid). Based on the results presented here, we would expect that the Bernardi et al. results would be off by at least a normalization factor, and indeed, an offset is seen in Figure 10 that is consistent with the fact that we required a negative correction to $\tau_{\text{eff}}(z)$ ($C < 0$) in order to correctly normalize our results. We note that there may be additional systematic errors related to the choice of analytic shape for the continuum, which Bernardi et al. parametrized as a power law plus three Gaussians to approximate the emission lines in the forest.

For the second, iterative approach used by Bernardi et al. (2003), it is not clear that their results

were converged, as claimed. They used power laws normalized at 1450 Å as initial estimates of the continua, and refined the continuum shapes based on the inferred $F(z)$ calculated at each iteration. In tests with our data, we found that after 3–4 iterations, as used by Bernardi et al., this procedure roughly recovered the continuum shape plotted in Figure 5. However, the normalization of the continuum, and hence $F(z)$, scaled directly with the normalization of the initial power-law continuum. Moreover, although the changes in the continuum were minor after a few iterations, the continua were not converged. After many iterations, the continua evolved to incorporate all of the relative differences between composites, producing a constant $F(z)$. This appears to be a generic result when the continua are allowed to converge separately for each composite. If the continua were held to be the same for each composite, then the shape of the continuum did converge, producing a relative $F(z)$ very similar to the one presented here. The overall normalization, however, did not converge. This simply demonstrates the fact that, without a completely unabsorbed composite to serve as a reference, the composites themselves do not contain enough information to fully constrain both $F(z)$ and the unabsorbed continuum without making other assumptions.

Dall’Aglio et al. (2009) fitted continua to 1733 individual quasar spectra by combining localized spline fits with a redshift-dependent global correction estimated from mock spectra. This approach is in principle similar to the one used by Faucher-Giguère et al. (2008); however, there are at least two significant differences. First, the use of low-resolution data means that the corrections must be substantially larger, and will be significant even down to $z \sim 2$. Additionally, the mock spectra used by Dall’Aglio et al. were created by randomly generating absorption lines from fixed number density, column density, and Doppler parameter distributions, rather than produced from cosmological simulations, as used by Faucher-Giguère et al.. The continuum corrections estimated from these mock spectra will therefore be subject to inaccuracies in the line parameter distributions, which were themselves estimated from high-resolution spectra that may have suffered from continuum fitting errors. This systematic effect offers a reasonable explanation for the relative tilt in the evolution of $\tau_{\text{eff}}(z)$ measured by Dall’Aglio et al. with our respect to our values. We note, however, that Dall’Aglio et al. find a smooth evolution with redshift, similar to our results, without any evidence for the feature at $z \sim 3.2$ reported by Bernardi et al. (2003).

McDonald et al. (2005) also found a relatively smooth evolution of $\tau_{\text{eff}}(z)$ with redshift in a sample of 3035 quasars. We include this work for completeness, although the authors stress that their method for estimating individual quasar continua in the Ly α forest using principle component analysis was not necessarily converged. They also required a normalization factor taken from other measurements, though this would not affect the overall shape of $\tau_{\text{eff}}(z)$. In general, their results reinforce the general conclusion that the mean transmitted flux measured from SDSS spectra evolves smoothly with redshift, consistent with the results presented here.

Finally, Pâris et al. (2011) recently applied PCA to measure the mean transmitted Ly α flux in the SDSS spectra of 2576 quasars. Their approach used principle components

derived from $z \sim 3$ quasar spectra to predict the unabsorbed continuum in the Ly α forest of individual objects based on the flux redward of Ly α . Overall, their values are generally consistent with our results, although very mild systematic differences can be seen in Figure 10. We caution that, despite the apparent agreement, a direct comparison with their results may not be entirely appropriate. Although Pâris et al. addressed the issue of optically thick systems by avoiding lines of sight containing DLAs, they did not correct for metal line contamination. Since they were measuring the mean flux directly from continuum-normalized spectra, such a correction would have been necessary to calculate the mean transmitted flux due to Ly α alone. We note that applying an adjustment of $\Delta\tau_{\text{eff}} \simeq -0.03$ for metals (see Section 2.2.1) would bring their $\tau_{\text{eff}}(z)$ values systematically below our one-sigma limits everywhere except in their three highest redshift bins.

5 SUMMARY

We have presented new measurements of the mean transmitted flux in the Ly α forest over $2.15 \leq z \leq 4.85$ made using composite quasar spectra. Our sample includes 6065 individual spectra from the SDSS DR7 quasar catalogue, which were combined into 26 composites with mean quasar redshifts $z_{\text{comp}} = 2.25$ to 5.08.

The fact that the quasar composites all have highly similar fluxes redward of Ly α suggests that the differences in the observed flux in the forest are driven by the evolution in the mean intergalactic Ly α opacity. The ratio of $F(z)$ at two different redshifts can therefore be measured by comparing the observed mean flux in two composites at the same rest wavelength. We used these ratios to directly constrain $F(z)$ as a fraction of $F(z = 2.15)$. The overall normalization for $F(z)$ was then determined using literature values made from high-resolution data at $z \leq 2.5$, where continuum uncertainties, although accounted for, are relatively minor. The statistical accuracy offered by the quasar spectra means that this normalization factor now dominates the uncertainty in $F(z)$ at $z \leq 3.75$.

Our main results give the mean transmitted flux associated with systems of column density $N_{\text{HI}} < 10^{19} \text{ cm}^{-2}$, and do not include the contribution from metal lines. We fit the mean flux using two functions. The first is a discrete $F(z)$ in bins of $\Delta z = 0.1$, while the second is a modified power law of the form $\tau_{\text{eff}}(z) = \tau_0 [(1+z)/(1+z_0)]^\beta + C$, where $\tau_{\text{eff}}(z) = -\ln F(z)$. Although previous works had generally parametrized $\tau_{\text{eff}}(z)$ as a pure power law ($C = 0$), we found that this was insufficient to describe the observed flux ratios at the required level of precision while simultaneously producing the correct normalization. We emphasize that the true functional form of $\tau_{\text{eff}}(z)$ is unknown, and that our parametrization is appropriate only in that it provides a reasonable match to the discrete $F(z)$ values. Caution should be used when extrapolating this (or any) fit to higher and lower redshifts. We also stress that, due to the fact that our results are generated from the ratios of fluxes at different redshifts, the parameters for both the discrete $F(z)$ and modified power law $\tau_{\text{eff}}(z)$ are correlated, and that the full correlation matrices should be considered when applying our results.

The mean transmitted flux is found to evolve smoothly with redshift, which is consistent with a gradual evolution of the H I ionization rate and gas temperature over $2 < z < 5$. We see no evidence of previously reported features at $z \simeq 3.2$, which appear to have been spurious. Our results are otherwise generally consistent with the most careful measurements made to date using high-resolution data (Faucher-Giguère et al. 2008), but offer a substantially greater level of precision and extend out to higher redshifts. Our measurements also offer an improvement over previous efforts using SDSS spectra, which appear generally to have been affected by systematic errors (although see Pâris et al. 2011).

This work demonstrates the significant advantages of using large data sets to measure the mean transmitted flux. In addition to the high degree of statistical accuracy, the use of composites makes it possible to avoid fitting continua at high redshifts, which is the primary source of systematic uncertainty in other approaches. These measurements should therefore serve as a reference for a variety of studies that make use of the Ly α forest.

ACKNOWLEDGEMENTS

The authors thank Claude-André Faucher-Giguère and Isabelle Pâris for their helpful comments on a first draft of this paper. GDB gratefully acknowledges support from the Kavli Foundation. PCH has received support from the STFC-funded Galaxy Formation and Evolution program at the Institute of Astronomy. GW and JXP have been supported through an NSF CAREER grant (AST-0548180) and by NSF grant AST-0908910.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is <http://www.sdss.org/>.

The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

REFERENCES

- Allen, J. T., Hewett, P. C., Maddox, N., Richards, G. T., & Belokurov, V. 2011, *MNRAS*, 410, 860
- Becker, G. D., Bolton, J. S., Haehnelt, M. G., & Sargent, W. L. W. 2011, *MNRAS*, 410, 1096
- Becker, G. D., Rauch, M., & Sargent, W. L. W. 2007, *ApJ*, 662, 72
- Bernardi, M., et al. 2003, *AJ*, 125, 32
- Bolton, J. S., Haehnelt, M. G., Viel, M., & Springel, V. 2005, *MNRAS*, 357, 1178
- Bolton, J. S., Oh, S. P., & Furlanetto, S. R. 2009, *MNRAS*, 396, 2405
- Dall’Aglio, A., Wisotzki, L., & Worseck, G. 2008, *A&A*, 491, 465
- . 2009, *arXiv:0906.1484*
- Dawson, K. S., et al. 2012, *arXiv:1208.0022*
- Fan, X., et al. 2006, *AJ*, 132, 117
- Faucher-Giguère, C.-A., Prochaska, J. X., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008, *ApJ*, 681, 831
- Hewett, P. C., & Wild, V. 2010, *MNRAS*, 405, 2302
- Kim, T.-S., Bolton, J. S., Viel, M., Haehnelt, M. G., & Carswell, R. F. 2007, *MNRAS*, 382, 1657
- Kirkman, D., et al. 2005, *MNRAS*, 360, 1373
- Lidz, A., Faucher-Giguère, C.-A., Dall’Aglio, A., McQuinn, M., Fechner, C., Zaldarriaga, M., Hernquist, L., & Dutta, S. 2010, *ApJ*, 718, 199
- McDonald, P., & Miralda-Escudé, J. 2001, *ApJ*, 549, L11
- McDonald, P., Miralda-Escudé, J., Rauch, M., Sargent, W. L. W., Barlow, T. A., Cen, R., & Ostriker, J. P. 2000, *ApJ*, 543, 1
- McDonald, P., et al. 2005, *ApJ*, 635, 761
- Pâris, I., et al. 2011, *A&A*, 530, 50
- Prochaska, J. X., O’Meara, J. M., & Worseck, G. 2010, *ApJ*, 718, 392
- Prochaska, J. X., & Wolfe, A. M. 2009, *ApJ*, 696, 1543
- Prochaska, J. X., Worseck, G., & O’Meara, J. M. 2009, *ApJ*, 705, L113
- Rauch, M., et al. 1997, *ApJ*, 489, 7
- Schaye, J., Aguirre, A., Kim, T.-S., Theuns, T., Rauch, M., & Sargent, W. L. W. 2003, *ApJ*, 596, 768
- Schaye, J., Theuns, T., Rauch, M., Efstathiou, G., & Sargent, W. L. W. 2000, *MNRAS*, 318, 817
- Schneider, D. P., et al. 2010, *AJ*, 139, 2360
- Songaila, A. 2004, *AJ*, 127, 2598
- Theuns, T., Bernardi, M., Frieman, J., Hewett, P., Schaye, J., Sheth, R. K., & SubbaRao, M. 2002, *ApJ*, 574, L111
- Vanden Berk, D. E. V., et al. 2001, *AJ*, 122, 549
- Viel, M., Becker, G. D., Bolton, J. S., Haehnelt, M. G., Rauch, M., & Sargent, W. L. W. 2008, *Physical Review Letters*, 100, 041304
- Wild, V., & Hewett, P. C. 2005, *MNRAS*, 358, 1083
- . 2010, *arXiv:1010.2500*
- Worseck, G., & Prochaska, J. X. 2011, *ApJ*, 728, 23
- York, D. G., et al. 2000, *AJ*, 120, 1579
- Zaldarriaga, M., Hui, L., & Tegmark, M. 2001, *ApJ*, 557, 519